

Avoiding Job Failure from Overfilling /PBS/spool

When your PBS job is running, its error and output files are kept in the `/PBS/spool` directory of the first node of your job. However, the space under `/PBS/spool` is limited, and when it fills up, any job that tries to write to `/PBS/spool` may die. This makes the node unusable by jobs until the spool directory is cleaned up manually.

To avoid this situation, PBS enforces a limit of 200 MB on the combined sizes of error and output files produced by a job. If a job exceeds that limit, PBS terminates the job.

To prevent this from happening to your job, do *not* write large amounts of content in the PBS output/error files. If your job needs to produce a lot of output, you can use the `qsub -kod` option to write the output directly to your final output file, bypassing the spool directory (see [Commonly Used qsub Options](#)). Until then, you can redirect standard out or standard error within your PBS script. Here are a few options to consider:

1. Redirect standard out and standard error to a single file:

```
(for csh)
mpiexec a.out >& output
(for bash)
mpiexec a.out > output 2>&1
```

2. Redirect standard out and standard error to separate files:

```
(for csh)
(mpiexec a.out > output) >& error
(for bash)
mpiexec a.out > output 2> error
```

3. Redirect only standard out to a file:

```
(for both csh and bash)
mpiexec a.out > output
```

The files "output" and "error" are created under your own directory and you can view the contents of these files while your job is still running.

If you are concerned that these two files could get clobbered in a second run of the script, you can create unique filenames for each run. For example, you can add the PBS JOBID to "output" using the following:

```
(for csh)
mpiexec a.out >& output.$PBS_JOBID
(for bash)
mpiexec a.out > output.$PBS_JOBID 2>&1
```

where `$PBS_JOBID` contains a number (jobid) and the name of the PBS server, such as `12345.pbspl1.nas.nasa.gov`.

If you just want to include the numeric part of the PBS JOBID, do the following:

```
(for csh)
set jobid=`echo $PBS_JOBID | awk -F . '{print $1}'`
mpiexec a.out >& output.$jobid

(for bash)
export jobid=`echo $PBS_JOBID | awk -F . '{print $1}'`
mpiexec a.out > output.$jobid 2>&1
```

If you used the `qsub -kod` option mentioned above, you will be able to see the contents of the output or error files directly from a front end as they grow. If you neither used `-k` nor redirected the output, you can still see the contents of your PBS output/error files before your job

completes by following these steps:

1. Find out the first node of your PBS job using `qstat -W o=+rank0`. In the following example, the output shows that the first node is `r162i0n14`:

```
%qstat -u your_username -W o=+rank0
JobID      User      Queue  Jobname  TSK Nds      wallt S      wallt  Eff Rank0
-----
868819.pbspl1 zsmith long   ABC      512 64 5d+00:00 R 3d+08:39 100% r162i0n14
```

2. Log into the first node and `cd` to `/PBS/spool` to find your PBS stderr/out file(s). You can view the contents of these files using `vi` or `view`.

```
%ssh r162i0n14
%cd /PBS/spool
%ls -lrt
-rw----- 1 zsmith a0800 49224236 Aug 2 19:33 868819.pbspl1.nas.nasa.gov.0U
-rw----- 1 zsmith a0800 1234236 Aug 2 19:33 868819.pbspl1.nas.nasa.gov.ER
```

Article ID: 183

Last updated: 04 Mar, 2020

Revision: 33

Running Jobs with PBS -> Optimizing/Troubleshooting -> Avoiding Job Failure from Overfilling /PBS/spool

<https://www.nas.nasa.gov/hecc/support/kb/entry/183/>